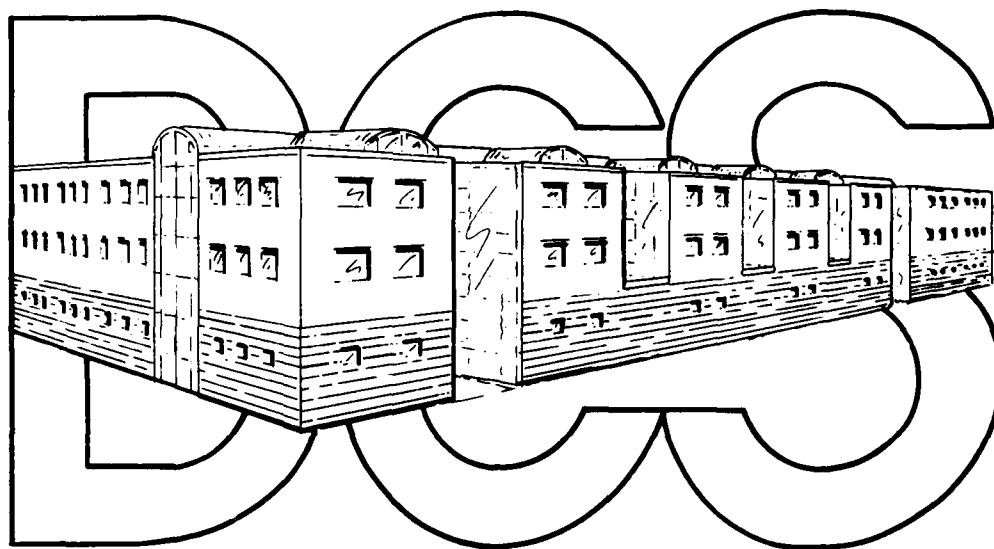COPY

DEPARTMENT OF COMPUTER SCIENCE

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

THE NEW ADDITION

REPORT NO. UIUCDCS-R-90-1577 UILU-ENG-90-1716

EXPLANATION-BASED LEARNING WITH PLAUSIBLE INFERENCING

by

Gerald DeJong

March 1990

DTIC
ELECTE
APR 2 5 1990
S D

90 04 23 117

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION		1b. RESTRICTIVE MARKINGS
Unclassified		None

2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	Approved for public release; Distribution unlimited

4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S)
UIUCDCS-R-90-1577	

6a. NAME OF PERFORMING ORGANIZATION	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION
Dept. of Computer Science University of Illinois		Office of Naval Research

6c. ADDRESS (City, State, and ZIP Code)	7b. ADDRESS (City, State, and ZIP Code)
1304 W. Springfield Urbana, Illinois 61801	800 N. Quincy Street Arlington, VA 22217-5000

8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
		N00014-86-K-0309

8c. ADDRESS (City, State, and ZIP Code)	10. SOURCE OF FUNDING NUMBERS			
	PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO	WORK UNIT ACCESSION NO.

11. TITLE (Include Security Classification)

Explanation-Based Learning with Plausible Inferencing

12. PERSONAL AUTHOR(S)
Gerald DeJong

13a. TYPE OF REPORT	13b. TIME COVERED		14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Technical	FROM _____	TO _____	1990 March	15

16. SUPPLEMENTARY NOTATION

17.	COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Explanation-Based Learning, Planning, Learning to Plan, Continuous Domains, Knowledge-Level Learning.
05		10	

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

This paper represents a synthesis of ideas from qualitative reasoning and explanation-based learning. Taken together they form a novel approach to planning that relies on *plausible inferencing* and applies to continuously varying rather than discrete world states. Interestingly, the frame problem skirted and the approach admits some forms of planning under uncertainty. Planning in a domain is very efficient, although learning about the domain can be time consuming. The approach possess a kind of natural reactivity.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT			21. ABSTRACT SECURITY CLASSIFICATION
☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS			Unclassified

22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL
Alan Meyrowitz	(202) 696-4302	ONR

REPORT NO. UIUCDCS-R-90-1577

EXPLANATION-BASED LEARNING WITH PLAUSIBLE INFERENCING

by
Gerald DeJong

March 1990

DEPARTMENT OF COMPUTER SCIENCE
1304 W. SPRINGFIELD AVENUE
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
URBANA, IL 61801

Explanation–Based Learning with Plausible Inferencing
Gerald DeJong
University of Illinois

This paper represents a synthesis of ideas from qualitative reasoning and explanation–based learning. Taken together they form a novel approach to planning that relies on *plausible inferencing* and applies to continuously varying rather than discrete world states. Interestingly, the frame problem skirted and the approach admits some forms of planning under uncertainty. Planning in a domain is very efficient, although learning about the domain can be time consuming. The approach possess a kind of natural reactivity.

1. Introduction

This paper investigates the application of Explanation–Based Learning (EBL) to planning tasks in continuous domains. To the extent that EBL has been applied to planning, it has been done so almost exclusively using some kind of of situation formalism whether the situation designator is explicit as in situation calculus or implicit as in a STRIPS formalism [Chien87, Mitchell85, Segre87, Simmons89] It is perhaps not surprising that the favored formalism for learning to plan is also the most common formalism for planning research. Situations are quiescent, permitting no change in truth–value of any formula describing the situation. Actions (or operator applications) are modeled as instantaneous mappings between situations. Situation formalisms offer a mechanism by which standard inferencing and theorem proving techniques can be applied to problem–solving and planning. Further, it supports the conventional EBL notion of an explanation as a truth–entailment proof, allowing conventional EBL to be performed in planning domains.

Unfortunately, situation formalisms saddle an EBL planning system with many undesirable characteristics as well. These characteristics which are sufficient to preclude most real–world applications. Chief among these, and the two that concern this paper, are 1) EBL problems inherent in the use of logical truth–entailment proofs as explanations and 2) the inadequacies of situation formalisms conveniently to handle temporally non–trivial states and actions.

There have been several temporal logics developed to address the second of these problems [Allen83, Dean83, McDermott82, Shoham86]. Their solutions are not conducive to EBL. The straight–forward use of any of these temporal logics yields monstrous proofs which result in awkward and expensive–to–match EBL concepts. Furthermore, use of such temporally logics additionally burden the implementer with the responsibility to define world mappings for combinations of overlapping actions. Finally, all are firmly cast in truth–entailment logic formalisms and therefore exhibit the inadequacies of logical proofs as explanations.

2. What Is Hard?

To illustrate the manifestations of these two problems for EBL consider what a human pupil might learn from the following two very brief scenario. The pupil is learning to fly an airplane. He is watching as the instructor approaches for a landing. The plane is on the correct glide slope to land but with an airspeed that is too fast:

> The instructor snuffs out his Benson and Hedges, adjusts his sunglasses, and gently closes the throttle while simultaneously easing the stick back which deflects the elevator up and increases the plane's angle of attack.

Why would this scenario be difficult for a current EBL system? First, there is no easy logical proof that the instructor's actions have the desired effect. The underlying aerodynamic theory necessary to support such a proof would be nearly impossible to give to the computer. Furthermore, most of the data necessary to complete such a proof is lacking. Missing information includes how much the fuel flow is

---

reduced by the observed throttle change, the effect this has on the energy output of the plane's engine, the efficiency of the propeller at the current RPMs, the local air density, the wing's lift coefficient, the aircraft's gross weight, air turbulence patterns between the plane and the runway, and much much more.

For current EBL systems, explanations are logical proofs. If there is no logical proof, there is no explanation, and no new concept can be learned. The identification of explanations with proofs is forced upon us by situation formalisms. Unfortunately, the number of real–world domains in which planning concepts can be supported by logical proofs is vanishingly small. Even for these few cases, McCarthy's frame problem and qualification problem [McCarthy69] impose staggering obstacles. If EBL is relegated to these degenerate domains one must question the significance of EBL for planning.

A second problem in the above scenario stems from the temporally interesting relation among the actions. Closing the throttle and moving the control stick must be done simultaneously and gradually. Yet, situation formalisms do not easily support temporal modeling of persistent actions. The standard response is simply to pretend that the actions are instantaneous in the domain theory while allowing them to persist in the real world. This trick cannot be used here. The two actions must be performed in a coordinated fashion. That is, the actions not only overlap temporally but must progress at the correct rate relative to each other. Such reasoning about gradual changes in world during the execution of operators is particularly difficult in situation formalisms.

The second scenario follows the first:

> When the plane has descended to several feet above the runway, the instructor reduces power, eases back further on the stick, and holds the controls motionless as the plane settles gently onto the runway.

Here we have another problem. In situation formalisms, all changes are due to actions; between actions a situation is quiescent and timeless. This is, of course, not the case in our world. In the above scenario there is no plausible action to attribute the effect of the plane's touchdown. It occurs in the middle of what should be a quiescent situation and yet there is an undeniable difference before and after touchdown (for example, the engine can be shut off with impunity after but not before, structural failure of the control surfaces would have qualitatively very different effects, etc.). Situation formalisms again provide us with no pleasant mechanism for modeling such world scenarios.

We must resort to extreme measures to preserve the situation formalism in the face of such problems. Typically, one would introduce some agentless action (that we might call "TOUCHDOWN") with appropriate preconditions and effects to map the "flying" situation to a "landed" situation in which other rules apply. For the first scenario one might introduce a set of cross–product operators which model as one unified whole the execution of several simultaneous operators (like closing the throttle and easing the stick back at a certain ratio).

While always possible, this approach does violence to the underlying situation. It postulates semantic items which have no correlate in the real world. In part this is only an aesthetic criticism: We have violated the underlying paradigm of designing AI systems. Before we projected some truth of the world onto a computational model which mirrored the world in all important ways while masking unimportant distinctions. Now we are suddenly forced to make distinctions which are not rooted in the world but exist purely for the sake of the computation.

Worse than that, this "fix" for situation formalisms places an unfair and under–constrained burden on the system implementor. It is he who must decide when to invent a new distinction and it is he who must craft that distinction to just make up for the deficiencies of his formalism. His job has clearly become much more difficult.

3. Plausible Explanations

A cornerstone of this research is the notion of a plausible explanation. With a conventional domain theory the explanation of a proposition is a logical proof of the proposition in terms of what is known. In a

theory of plausible inference a proof is an educated, somewhat abstract guess at why the proposition is likely to be true given what is believed. For example, one might plausibly reason that since it is autumn in Central Illinois, tomorrow will be a windy day. This illustrates the two hallmarks of our plausible inferences: First, they are not certain. It is entirely possible that tomorrow will not, in fact, be windy in Central Illinois. Second, plausible inferences are often abstract. It is not plausible to conclude that the winds will be out of the north northwest at 22 mph. To be an acceptable rule the characterization of the wind must be much more abstract. Often plausible conclusions are under–specified in this way.

We will refer to these abstract, under–determined proof–like objects as a "plausible explanation" for a conclusion according to a plausible theory. It is important not to abuse the term "proof" which is closely linked to truth entailment deductions. While proofs figure prominently in some approaches to FBL [Mitchell86] we will avoid that term. Explanations play the role of logical proofs for a conventional EBL system. As in a conventional EBL system, the plausible explanation is generalized to form a new concept. There is, however, no guarantee that the resulting concept will be correct. In this way plausible explanations are more abductive than deductive.

A plausible explanation also has the dual properties of uncertainty and imprecision. Perhaps surprisingly, these properties when applied to general concepts rather than specific instances can be advantageous to an explanation–based learning system.

The "plausible" character of the explanation stems from the domain theory. We will look at the characteristics of uncertainty and imprecision in individual rules and then see how these characteristics can be beneficial for generalized planning concepts.

The constituents of a plausible theory may be syntactically indistinguishable from a conventional domain theory (e.g., horn clauses, production rules, schemata, etc.) but they have a rather different semantics. Consider a simple implicature: "A => B". The conventional semantic interpretation of this rule is straightforward. If it is to be used as a plausible rule, its meaning changes significantly.

3.1. Uncertainty

A plausible rule may not always reflect reality. There may be conditions under which "A" is satisfied but "B" is not true. The equivalent rule under a standard semantics would be of the form "C & A => B". Here C represents a specification of the context in which the plausible rule is guaranteed. C specifies the implicit assumptions built into the plausible rule "A => B".

To be a useful rule to the plausible inference system, the conditions that make C false should be, for the most part, infrequent or otherwise uninteresting. This idea is similar in spirit to Winston's censors [Winston86] except that Winston makes the censors explicit. It is essential that the precise conditions of applicability of a plausible rule remain implicit.

Much of the power of the current approach is traceable to the fact that no attempt is made to specify the context conditions of a domain rule (such as "C" in the above implication). Such context conditions must not be represented or directly reasoned about. It is also important that the uncertainty in a rule not reflect uncertainty in the domain itself. A rule capturing the fact that when two dice are thrown they come up seven more than any other number, for example, is a rule that reflects domain uncertainty. We have not investigated EBL applied to rules with this sort of uncertainty.

3.2. Imprecision

The constraint specified by a plausible rule is often a generality or abstraction. This is because the expressions that compose plausible inference rules themselves often refer to abstractions. In a plausible rule "A => B" the symbols A and B will seldom refer to precise world objects. Consider a plausible rule in the domain of driving an automobile which states "pressing down on the accelerator makes the car go faster". The antecedent "pressing down on the accelerator" is not a precisely defined event description, nor is the consequent well defined (the car going faster). Both describe rather broad classes of events. The rule does not say precisely what degree of acceleration should be expected in a particular car from a certain pressing action. It only says that some amount of acceleration will likely be found if some amount of

pushing on the gas pedal is performed.

Such imprecision is the source of a kind of generality. This imprecision–generality, like the sort of uncertainty described above, can be turned to the advantage of an explanation–based learning system. But again, the plausible theory must have the right sort of imprecision. What sorts of imprecision are useful? One of the essential characteristics is that the imprecision be contained wholly in the rule and not reflect a fundamental domain fuzzyness. This is the case with the example cited above. The fuzzy description of a car going faster is not a fuzzyness of the domain. Cars cannot "go faster" abstractly. Real world cars change from one specific velocity to another. Another essential characteristic for advantageous imprecisions is *continuity*. By this we mean that in the cases that the rule applies, the two actual real world consequents can be made arbitrarily close by insuring that the differences between the antecedents are sufficiently small if all other things are equal. This condition is met by our acceleration rule: If one pushed a *slightly* harder on the accelerator in one instance than another, the resulting velocity would also be only slightly different, assuming the brake is not depressed in one instance and not the other, etc. Note that this notion of plausible inference is significantly different than [Collins86].

4. Plausible EBL

We adopt the standard learning–apprentice paradigm as illustrated in [Segre87] and [Mitchell85]. The system is given a set of general goal specifications. These goals are beyond the systems initial abilities and it must learn how to achieve the general goal classes by observation of an expert. It monitors the behavior of the expert and observes how the world changes. When it detects that the expert has achieved an example of one of its general goals an explanation is constructed. Once generalized it forms the basis of a new planning concept.

Certain items in the world are directly observable by the system. We require the abstract goal set to be specified entirely with observables to facilitate the detection of achieved goals. In addition, the values of some variables are directly controllable. Performing an action is just the specification of how one or more of these controllable parameters are changed. There is no assumption that the changes are instantaneous or that other items are quiescent between actions.

An observation is the specification of values for all observables over some time interval. Plausible EBL involves generalizing the explanation of an observation, just as in conventional EBL. However, this alone is insufficient to generate a useful planning concept. There are two reasons: 1) the explanation might be wrong and 2) the planning concept resulting from the generalization may be too abstract and imprecise to be applicable.

The imprecision of parts composing the new concept first appear to compromise the concept's applicability. One seldom wants to achieve the goal of making a car go some (unpredictable) amount faster; rather one has specific goals such as accelerating my 1981 Toyota to 55 mph. In fact if the domain theory contains only the right sort of imprecision, that imprecision can be an asset. The plausible explanation, while imprecise, is deterministic. Achieving the preconditions in the same way on repeated applications will result in the same consequents. This is similar to Russell's notion of *determinations* [Davies87] generalized to continuous domains. The concept's imprecision precludes deducing what those consequents will be. But a mechanism external to the generalized explanation can remember from one application to the next how to pair specific antecedents to specific consequents and how to hypothesize precise consequents for unseen antecedents provided they follow the pattern set by the known antecedents. With the addition of a simple interpolation function, they can be used to form explanation–based concepts that support planning. The interpolation function relates precise values of antecedents and consequents. Every time a plausible EBL concept is actually applied, the antecedent/observed–consequent pair is asserted to the concept's interpolation function. If either antecedents or consequents are not known, a precise (though possibly somewhat incorrect) value for the other can be provided by interpolating among the known points. The quality of the interpolated value depends on the number and distance of nearby already–known points, and on how smooth the antecedent/consequent function is in the vicinity of the new point. With sufficient experience the new concept's interpolation function can guarantee arbitrarily small errors. [Haggerty72].

But what about uncertainty? The previous discussion is predicated upon the assumption that the explanation is faithful to the world. How can this be known? The answer is that it cannot be known for certain. However, there are a number of reasons that this problem is not as devastating as it may appear. First, the new concept is a plausible concept. That is, it is valid with respect to the plausible theory. If our plausible theory is indeed plausible, it must at a minimum guarantee that WFFs with plausible derivations are more likely to describe the world accurately than arbitrarily constructed WFFs. Indeed, we may take this to be an informal definition of "plausibility": a theory is a plausible theory iff the conditional probability that a concept is faithful to the world given that there is a derivation in the theory is greater than its *a priori* probability but less than one. This definition also supports a reasonable "plausibility" ordering relation among theories that cover the same concepts. Since the new concept is derivable from the theory there is some justification in believing the concept describes the world. The strength of this justification is related to the plausibility of the theory.

Second, there is information inherent in the fact that the training example itself occurred in the world. Consider the set of WFFs derivable from the plausible theory. Some (perhaps many) are incompatible with the world. For these, there can be no confirming observation. Insisting that at least one real-world example be found eliminates WFFs which are uniformly unfaithful to the world. The conditional probability that a plausible WFF is faithful to the world given that a world example has been seen, is higher than the *a priori* probability that a plausible WFF is faithful to the world. Plausible EBL must not go off generalizing randomly generated deductions. New concepts are formed only in the context of a training example. The existence of the training example itself adds credibility to the faithfulness of the plausible explanation and, therefore, to the new generalized concept. Relying on observations in the world to contribute to the plausibility of an explanation considerably elevates the role of the training example in plausible explanation–based systems as compared with conventional explanation–based generalization systems (e.g., [Mitchell83, Mitchell86, Mooney88]. Third and finally, concepts formed from the generalization of unfaithful explanations can in all interesting cases be detected as incorrect while the system uses the concept. This is an extra service provided by the interpolation function. The interpolation function remembers relations among quantities as they were actually observed to exist in the real world. The selection of which variables to include in the interpolation is provided by the generalized plausible explanation. The explanation contends that for world situations of interest there is a systematic relation among the components of its antecedents and consequents. If this is not true then the interpolation function will sooner or later be required to include a point that is inconsistent with existing points. This is a clear indication that the relation among the quantities is not systematic, as required by the explanation. Thus, the explanation is not faithful to the world.

The plausible theory may impose smoothness constraints in addition to mere continuity constraints. This is highly desirable. In addition to the requirements imposed on the sorts of uncertainties and imprecisions, theories are most conducive to plausible explanation–based learning if they enforce some smoothness constraint on the interpolation function. The smoother the function can be guaranteed to be, the more constrained are the points, and fewer examples will be needed for the interpolation function to detect an inconsistency. An related benefit from a smooth concept interpolation function is that each asserted point has greater predictive power; smooth functions are more easily interpolated so the interpolation error decreases more rapidly.

The plausible EBL algorithm:

Given a set G of abstract goals of interest
      a plausible domain theory including the specification of observable
         and controllable quantities
      an interpolation strategy consistent with the domain theory
      a noise threshold for each observable
1) Monitor the actions of an expert for the achievement of an element, GOAL,
      of the set G. (Alternatively, select an element and ask the expert to
      achieve it.) Collect the monitored world observables, the expert actions,
      and the instance of GOAL achieved. Call the collection EXAMPLE.
2) Construct EXPLANATION, a new plausible explanation for EXAMPLE using
      the domain theory. If there are no remaining explanations for EXAMPLE
      then give up; the plausible domain theory is insufficient to support the
      observation. Otherwise continue.
3) Generalize EXPLANATION. Compute a description of the part of GOAL that
      can be achieved by the generalization of EXAMPLE by intersecting GOAL
      and the EBL generalized version of EXPLANATION's
      conclusion. Call this CONCEPT-GOAL. For CONCEPT-GOAL compute
      the generalized preconditions. (This may be done efficiently using
      EGGS [Mooney86] or EBG's limited goal regression
      [Mitchell86]. Call this CONCEPT-PRECONDITIONS. Collect
      the relevant expert actions of EXPLANATION with generalized parameters
      into a list called ACTIONS. Construct an empty interpolation function,
      IF, for the parameters of ACTIONS and CONCEPT-GOAL. Call the combination
      of CONCEPT-GOAL, CONCEPT-PRECONDITIONS, ACTIONS, and IF "CONCEPT". 4)
   Assert the observed points from EXAMPLE into IF. 5) Use the new concept for planning:
    A) When a new goal is given to the system which unifies with CONCEPT-GOAL
      and CONCEPT-PRECONDITIONS can
      be met, IF to compute specific
      parameters for elements of ACTIONS. Execute the specific actions
      in the real world, and observe the results.
    B) Was the goal achieved to within the noise threshold for each observable?
      i) If yes, go to (5)
      ii) If no, assert the observed result IF.
         If IF can successfully integrate the new values go to (5).
      iii) If the new values cause a contradiction within IF,
         throw out CONCEPT and go to (2).

If the domain theory is sufficient to represent an example, and if explanations are recursively enumerable, it can be shown that the algorithm will converge to an adequate planning concept that includes the example.

It is advantageous to be able to generate explanations ordered by plausibility with the most plausible first, although this is not required for convergence. Another interesting point is that throwing the concept out in step (iii) is not necessary, although it makes the analysis much easier. Instead, if the concept has proved to be useful one could try to narrow its preconditions empirically to avoid function values in the neighborhood of the inconsistent point or simply keep the concept and tolerate a few errors in its application.

5. Qualitative Reasoning

Are there indeed domain theories that satisfy our criteria for plausible theories? Yes, many of the theories produced in the AI area of qualitative reasoning have just the characteristics needed. For the system we adapt a simplified version of Forbus' Qualitative Process Theory [Forbus84].

Domain knowledge is coded in the form of *processes*. Each process has preconditions and a body. The body specifies a set of plausible constraints among quantities. Over intervals in which the preconditions of a process are met, the process becomes active and its plausible constraints are available to the inferencer. Many processes may be active at once.

5.1. Types *Quantities* are world sate variables that can take on numeric values. While they take on numeric values, their values are reasoned about in the qualitative model in purely qualitative terms. All changes in a quantity's value are continuous.

Quantities are of three types: *observable, non-observable,* or *constant.* Constant quantities take on unchanging values. Whether a quantity is observable or non-observable depends on whether or not the system has direct access to the values taken on by the quantity. Observable quantities may be either *parameters* or *internal quantities.* Parameters are set by the environment and must be input to the system. The value of internal quantities are determined by the laws of nature from the values of parameters.

Parameters can be either *controllable* or *non-controllable* depending on whether or not the quantity's values are directly manipulable. Controllable quantities are things like the position of a radio's volume knob or the setting of a thermostat. Non-controllable parameters include things like the amplitude of sound waves from a radio's speaker or the density of air around an aircraft. These may be monitored but they can be changed only indirectly through manipulation of controllable parameters (if indeed they are changeable at all). Non-observables cannot be parameters in the current system. The system knows the world only through observable quantities. The world may be influenced only through controllable parameters.

5.2. Qualitative Predicates and Proportionalities

The qualitative descriptions for a quantity are: INCREASING, DECREASING, or CONSTANT, and they may be GREATER-THAN, LESS-THAN, or EQUAL to other quantities.

All plausible constraints are in the form of qualitative proportionalities which are binary relations over quantities. At any point the available qualitative proportionalities are the union of the bodies of active processes. Qualitative proportionalities may be positive or negative. The positive qualitative proportionality "Q+" is taken to mean that if one of its quantity arguments changes, the other will change in a like manner, all other things being equal. (Q+ A B I) represents that over the interval I, quantities A and B are positively qualitatively proportional. If A is known to be increasing on a sub-interval of I then B may plausibly be inferred to be increasing on that same sub-interval also. Likewise if A is decreasing on a sub-interval B may be also decreasing on that sub-interval. The negative qualitative proportionality "Q-" is similar except that from (Q- A B I) and knowing A is increasing it may plausibly be inferred that B is decreasing.

This formalism for plausible theories supports a necessary (though possibly surprising) property: inconsistent conclusions are allowed. There may be derivations both for A INCREASING over an interval and for A DECREASING over the same interval.

6. An Example

The following example demonstrates an implemented system acquiring a new planning concept in the domain of driving standard transmission automobiles. The system is implemented partially in NSW PROLOG and partially in LUCID CommonLisp. It runs on an IBM RT and is currently being re-written entirely in common lisp.

The general goal to be achieved is to cause the car to increase its speed. Included in this general specification is the possibility that the car is starting up from a dead stop (although with the engine idling). The solution often requires the execution of two actions at once: both gradually letting out the clutch while increasing the position of the gas pedal in a coordinated way. As there is only one interesting qualitative situation for this necessarily simple problem, time intervals are left out of the following discussion.

The only relevant qualitative process is:

PC: REVS > MIN-REVS
            REVS < MAX-REVS
            CLUTCH > FRICTION-POINT

BODY: (Q+ GAS-FLOW GAS)   (Q- REVS ENGAGEMENT)
            (Q+ REVS GAS-FLOW)   (Q- SPEED GRADE)
            (Q+ SPEED REVS)   (Q- REVS GRADE)
            (Q+ ENGAGEMENT CLUTCH) (Q+ TEMP SPEED)
            (Q+ SPEED ENGAGEMENT)   (Q+ REVS TEMP)

GAS and CLUTCH represent the position of the gas and clutch pedals respectively. Both are controllable. The zero position of GAS is fully up, the zero position for CLUTCH is fully depressed. ENGAGEMENT is the percentage of power transmitted to the wheels through the clutch. REVS is the rotational speed of the engine. SPEED is the speed of the car. GRADE is the hill gradient, a non-controllable parameter. TEMP is the engine temperature. GAS, CLUTCH, REVS, SPEED, TEMP, and GRADE are observable.

The process says that while the engine revolutions per minute is above a minimum threshold (so that the engine does not die), while it is below a maximum threshold (that would cause immediate and irreparable damage to the engine), and while the clutch is at least partially engaged, a number of qualitative proportionalities hold that capture roughly the following intuitions about automobiles: increasing GAS makes REVS go up and SPEED go up; letting out the clutch makes SPEED go up and REVS go down; a steeper grade causes SPEED and REVS both to go down; the engine heats up at higher speeds and allows more efficient combustion (REVS go up).

It was mentioned earlier that explanations should be generated in a most-plausible-first order. Such an order requires some kind of quantitative measure of plausibility. To simplify the system we assume that the plausibility of an explanation is approximated by the simplicity simplicity of the explanation. This is equivalent to assuming all rules have the same, independent plausibility. Thus, an explanation requiring 3 plausible rule applications will be generated before one with 4 or 5 rules applications.

The goal given to the system is (INCREASING SPEED). Initially, the system has no planning concept that is relevant to this goal. The system might have been designed to perform some kind of search through its rules to find a plausible way to increase SPEED. However, because the space is potentially large and filled with specious concepts it is extremely unlikely that any such strictly analytical approach would yield a solution. Instead, the system waits for an expert to provide it with an example. The expert's actions provide the following observation:
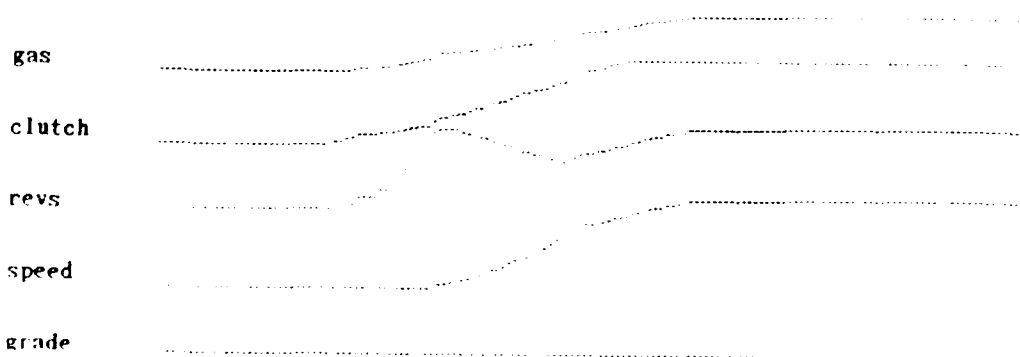

Figure 1:   Observed Expert Behavior

The first explanation constructed is:

```
(EXPLANATION (INCREASING SPEED)
        ((Q+ SPEED REVS)
         (Q+ REVS GAS-FLOW)
         (Q+ GAS-FLOW GAS)
         (INCREASING GAS)))
```

A quantitative interpolation function is created for the interval of applicability and the observed numerical values for SPEED, and GAS are asserted. The variables in interpolation function is the the goal quantity added to the set of all of the parameters mentioned in the explanation. These are SPEED and GAS. Another acceleration problem is given to the system. It is to accelerate from a slow speed with the clutch already partially engaged. The system selects the newly-constructed planning concept with the following r sults:
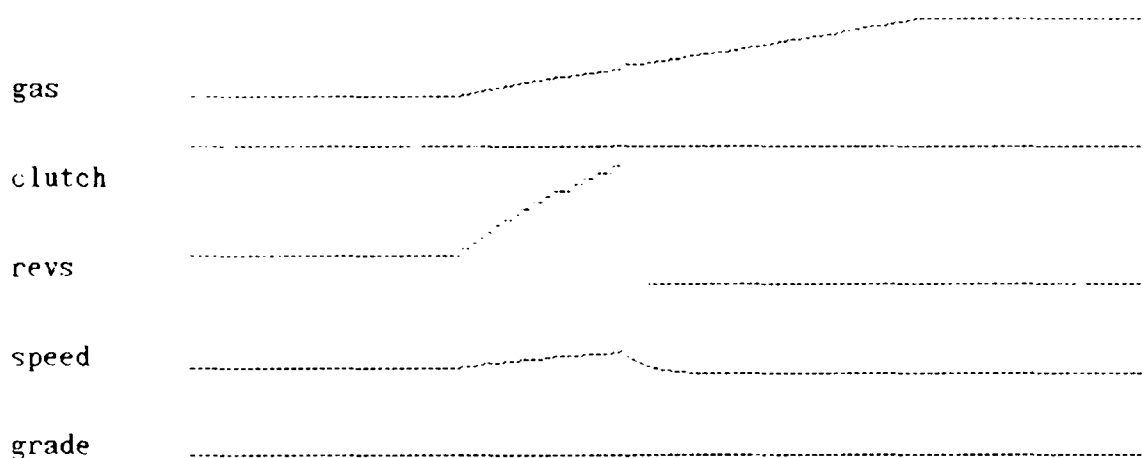
gas

clutch

revs

speed

grade

Figure 2:   System Behavior Using Gas-Only Concept

The explanation is, in fact, not the right one. Its veracity depends on implicit conditions which are not met in the example. The plan results in over-revving the engine which suffers irreparable damage and stops working. The observation is inconsistent with the qualitative explanation. The system searches for a plausible explanation for the the qualitative discrepancy (SPEED decreasing in spite of GAS increasing), a d finds the possibility that a precondition (REV < MAX-REV) was plausibly violated. The concept is discarded and another explanation is generated for the original expert observation but with the added requirement that REV be controlled to be less than MAX-REV. The following plausible explanation is constructed:

```
(EXPLANATION (AND (INCREASING SPEED) (REV < MAX-REV))
        ((Q+ SPEED ENGAGEMENT)
         (Q+ ENGAGEMENT CLUTCH)
         (INCREASING CLUTCH)
         (Q- REVS ENGAGEMENT)))
```

In planning with this concept the car is stalled. The concept fails less spectacularly, but it also yields real world results that are inconsistent with the explanation. Finally, a more adequate qualitative explanation is generated:

```
(EXPLANATION (AND (INCREASING SPEED) (< REVS MAX-REVS) (> REVS MIN-REVS))
              ((Q+ SPEED REVS)
               (Q+ REVS GAS-FLOW)
               (Q+ GAS-FLOW GAS)
               (INCREASING GAS)
               (Q- REVS ENGAGEMENT)
               (Q+ SPEED ENGAGEMENT)
               (Q+ ENGAGEMENT CLUTCH)
               (INCREASING CLUTCH)))
```

This explanation allows the SPEED to be unambiguously increasing while plausibly the REVS are managed within the bounds needed. Notice that the explanation is not, in fact, correct. It misses the effect of GRADE upon SPEED. If a world situation depended on the missing information, this concept also would be eliminated in favor of a more faithful explanation. Nonetheless, this time the quantitative interpolation function includes both CLUTCH and GAS as controllables. The system generates the following adequate solution to a posed problem:
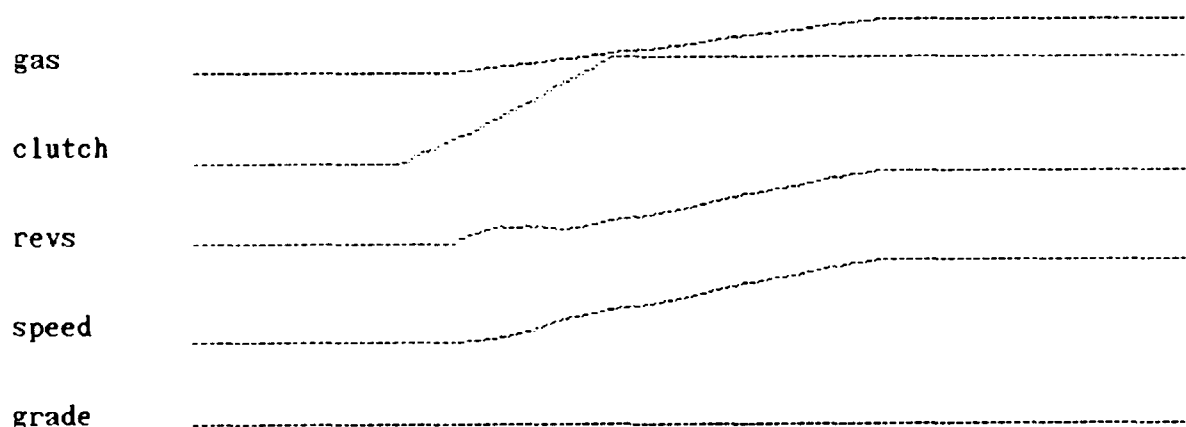
gas

clutch

revs

speed

grade

Figure 3:   System Behavior Using Gas-Clutch Concept

6.1. Conclusions

The paper describes a system that learns to plan in a world made up of continuously varying quantities. The system demonstrates a symbolic (that is non-connectionist) alternative to solving some of the problems underlying situation formalisms and other traditional logic-oriented formalisms. It shows how an unsound inference procedure can not only be useful but can skirt some troublesome problems faced by standard logical approaches to planning and action [Georgeff aaai86, Pednault aaai88]. The frame problem [McCarthy69], for example, does not rear its ugly head in planning with plausible explanation-based learning. The reason is that no attempt is made to guarantee a perfect plan the first time. Perfection comes with practice and use. The more a plausible concept is used, the more observed quantiative relations will be asserted to the interpolation function, and the smaller will be the error between the interpolated values and the values of the ideal function. Furthermore, the more often a concept demonstrates that a world situation does not contradict any of its myriad of implicit assumptions, the less likely it is to fail on an arbitrary future world situation.

Plausible explanation-based learning systems can learn at the knowledge level [Dietterich86, Newell81]. This is again traceable to the implicit assumptions. Every inferential step performed in the construction of an explanation has a built in "inductive leap of faith".

REFERENCES

[Allen83]   J. F. Allen, "Maintaining Knowledge about Temporal Intervals," *Communications of the Association for Computing Machinery 26*, 11 (November 1983), pp. 832–843.

[Chien87]   S. A. Chien, "Simplifications in Temporal Persistence: An Approach to the Intractable Domain Theory Problem in Explanation–Based Learning," M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, August 1987. (Also appears as UILU–ENG–87–2255, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

[Collins86]   A. Collins and R. Michalski, "The Logic of Plausible Reasoning: A Core Theory," UILU–ENG–86–1775, ISG, University of Illinois at Urbana–Champaign, 1986.

[Davies87]   T. R. Davies and S. J. Russell, "A Logical Approach to Reasoning by Analogy," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy, August 1987, pp. 264–270.

[Dean83]   T. Dean, "Time Map Maintenance," Technical Report 289, Yale University, New Haven, CT, October 1983.

[Dietterich86] T. G. Dietterich, "Learning at the Knowledge Level," *Machine Learning 1*, 3 (1986), pp. 287–316.

[Forbus84]   K. D. Forbus, "Qualitative Process Theory," *Artificial Intelligence 24*, (1984), pp. 85–168.

[Haggerty72] G. Haggerty, *Elementary Numerical Analysis with Programming*, Allyn and Bacon, Boston, 1972.

[McCarthy69] J. McCarthy and P. J. Hayes, "Some Philosophical Problems from the Standpoint of Artificial Intelligence," in *Machine Intelligence 4*, B. Meltzer and D. Michie (ed.), Edinburgh University Press, Edinburgh, Scotland, 1969.

[McDermott82] D. McDermott, "A Temporal Logic for Reasoning About Processes and Plans," *Cognitive Science 6*, 2 (1982), pp. 101–155.

[Mitchell83] T. M. Mitchell, "Learning and Problem Solving," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, West Germany, August 1983, pp. 1139–1151.

[Mitchell85] T. M. Mitchell, S. Mahadevan and L. I. Steinberg, "LEAP: A Learning Apprentice for VLSI Design," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 1985, pp. 573–580.

[Mitchell86] T. M. Mitchell, R. Keller and S. Kedar-Cabelli, "Explanation–Based Generalization: A Unifying View," *Machine Learning 1*, 1 (January 1986), pp. 47–80.

[Mooney86]   R. J. Mooney and S. W. Bennett, "A Domain Independent Explanation–Based Generalizer," *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, August 1986, pp. 551–555. (Also appears as Technical Report UILU–ENG–86–2216, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

[Mooney88]   R. J. Mooney, "A General Explanation–Based Learning Mechanism and its Application to Narrative Understanding," Ph.D. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, January 1988. (Also appears as UILU–ENG–87–2269, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

[Newell81]   A. Newell, "The Knowledge Level," *Artificial Intelligence Magazine 2*, (1981), pp. 1–20.

[Segre87]   A. M. Segre, "Explanation–Based Learning of Generalized Robot Assembly Tasks," Ph.D. Thesis, Department of Electrical and Computer Engineering, University of Illinois, Urbana, IL, January 1987. (Also appears as UILU–ENG–87–2208, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

[Shoham86]   Y. Shoham, "Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence," PhD. Thesis, Yale University, Dept. of Computer Science, New Haven, CT, 1986.

[Simmons89]   L. J. L. T. M. M. A. P. R. Simmons, "A Case Study in Robot Exploration," CMU-RI-Technical Report-89-1, Computer Science Department, Carnegie Mellon University, January 1989.

[Winston86]   P. H. Winston, "Learning by Augmenting Rules and Accumulating Censors," in *Machine Learning: An Artificial Intelligence Approach, Vol. II*, T. M. Mitchell, J. G. Carbonell and R. S. Michalski (ed.), Morgan Kaufmann, Los Altos, CA, 1986, pp. 45-61.

BIBLIOGRAPHIC DATA SHEET	1. Report No. UIUCDCS-R-90-1577	2.	3. Recipient's Accession No.
4. Title and Subtitle			5. Report Date March 1990
EXPLANATION-BASED LEARNING WITH PLAUSIBLE INFERENCING			6.
7. Author(s) Gerald DeJong			8. Performing Organization Rept. No. R-90-1577
9. Performing Organization Name and Address			10. Project/Task/Work Unit No.
Dept. of Computer Science 1304 W. Springfield Ave. Urbana, IL 61801			11. Contract/Grant No. N00014-86-K-0309
12. Sponsoring Organization Name and Address Office of Naval Research			13. Type of Report & Period Covered
			14.

15. Supplementary Notes

16. Abstracts

This paper represents a synthesis of ideas from qualitative reasoning and explanation-based learning. Taken together they form a novel approach to planning that relies on *plausible inferencing* and applies to continuously varying rather than discrete world states. Interestingly, the frame problem skirted and the approach admits some forms of planning under uncertainty. Planning in a domain is very efficient, although learning about the domain can be time consuming. The approach possess a kind of natural reactivity.

17. Key Words and Document Analysis. 17a. Descriptors

Explanation-Based Learning
Planning
Learning to Plan
Continuous Domains
Knowledge-Level Learning

17b. Identifiers/Open-Ended Terms

17c. COSATI Field/Group

18. Availability Statement	19. Security Class (This Report) UNCLASSIFIED	21. No. of Pages 15
unlimited	20. Security Class (This Page) UNCLASSIFIED	22. Price

FORM NTIS-35 (10-70)   USCOMM-DC 40329-P71